

Market Mechanism Design for Profitable On-Demand Transport Services

Malcolm Egan^{1,*} and Michal Jakob¹

¹ *Agent Technology Center, Faculty of Electrical Engineering, Czech Technical University in Prague, Czech Republic.*

Email: {malcolm.egan,michal.jakob}@agents.fel.cvut.cz

** Corresponding author.*

Abstract

On-demand transport services in the form of dial-a-ride and taxis are crucial parts of the transport infrastructure in all major cities. However, not all on-demand transport services are equal. In particular, not-for-profit dial-a-ride services with coordinated drivers significantly differ from profit-motivated taxi services with uncoordinated drivers. As such, there are two key threads of research for efficient scheduling, routing, and pricing for passengers: dial-a-ride services (first thread); and taxi services (second thread). Unfortunately, there has been only limited development of algorithms for joint optimization of scheduling, routing, and pricing; largely due to the widespread assumption of fixed pricing. In this paper, we introduce another thread: profit-motivated on-demand transport services with coordinated drivers. To maximize provider profits and the efficiency of the service, we propose a new market mechanism for this new thread of on-demand transport services, where passengers negotiate with the service provider. In contrast to previous work, our mechanism jointly optimizes scheduling, routing, and pricing. Ultimately, we demonstrate that our approach can lead to higher profits, compared with standard fixed price approaches, while maintaining comparable efficiency.

Keywords: on-demand transport; market mechanism; pricing

I. INTRODUCTION

Call a taxi in any major city and it will often arrive within minutes. Despite the success and widespread use of taxis and other on-demand transport services, there is room for improvement: higher profits; reduced prices; and even lower waiting times, targeted at passengers that have the highest demand for the service. Fortunately, the ubiquity of mobile internet and secure online financial transactions has opened the way for highly efficient on-demand transport; able to ensure that the right driver transports the right passenger at the right time.

At present, efficiently operating a fleet of on-demand vehicles remains difficult. The cause is the coupling between three key sub-problems: which passengers should be serviced by each vehicle (*routing*); what time should each passenger be picked up (*scheduling*); and how much each passenger should be charged (*pricing*). In effective on-demand transport services, the sub-problems should be solved jointly, not be decoupled.

Unfortunately, there are few computational techniques to solve the three on-demand sub-problems jointly; despite over three decades of research on related vehicle routing problems. While this might seem surprising, there is good reason: there are in fact two distinct research threads—each addressing a different niche.

The first thread is dial-a-ride services—targeted at the niche of elderly and disabled transport, as well as in low demand regions. Vehicles in fleets offering this type of service are coordinated by a single provider. Each vehicle collects passengers within requested pick-up time intervals and drops each passenger off before a requested drop-off time. Dial-a-ride services are typically heavily subsidized by governments, due to the important role they play for vulnerable members of the community; for instance, by taking an elderly woman to the hospital for a check-up. These subsidies have an important repercussion: dial-a-ride services are not-for-profit. As such, the aim of providers is to minimize costs; as opposed to maximizing profits. This means that the pricing sub-problem is not considered, and the most popular formalization known as the dial-a-ride problem focuses on the routing and scheduling sub-problems. Both centralized (see [1] for an extensive survey) and decentralized agent-based [2], [3] approaches have been proposed.

The second thread is taxi and private hire services. Taxi services contrast with dial-a-ride

services in two key ways: they are profit-motivated; and vehicles are not heavily coordinated by a single service provider. This is because the drivers are self-interested and unable to easily determine the current locations and destinations of other nearby taxis. While taxi services are profit-motivated, the most common pricing strategy is to use a fixed price-rate; i.e., the price scales as a—usually linear—function of the distance (see [4] for pricing with a non-linear function of distance). As such, the price-rate does not factor in the number and demand¹ of passengers that have ordered a ride. Moreover, the fixed price-rate means that the pricing sub-problem is decoupled from routing and scheduling; instead, the focus is on reducing waiting times and travel distance (and hence reducing costs) [5]–[9]. We note that the approach in [10] does optimize pricing for single taxi operation via dynamic programming; however, scheduling and routing for a fleet of taxis is not considered.

In essence, algorithms for the joint solution of the three on-demand sub-problems have not been developed in either of these well-established threads of on-demand transport research. As such, in this paper, we introduce a new thread—made practical with recent technological developments—where joint solution is feasible.

A. A New Thread

The aim of the new thread of on-demand transport services is profitability and efficiency when drivers are coordinated (as in dial-a-ride services) and providers are profit-motivated (as in taxi services). It is worth noting that this category only recently arose as more drivers and passengers have adopted internet-enabled smartphones, allowing for providers to easily keep track of demand statistics and for passengers to directly pay the provider (instead of the driver).

Looking from the computational perspective, standard algorithms for dial-a-ride or taxi services cannot be directly applied. This is due to the fact that the first two threads decouple pricing from routing and scheduling. As such, to solve the three sub-problems jointly, new computational techniques are required.

¹The notion of demand here can be easily misconstrued. We mean demand relative to factors such as cost of living and time of day. We do not mean in terms of need; in particular, we believe that pricing should not be manipulated in emergency situations.

In this paper, we propose a new market-based approach for the new thread of on-demand transport services. In fact, we jointly solve the three sub-problems: routing; scheduling; and pricing. There are two fundamental aspects of our approach: a new passenger model; and a new market mechanism. The passenger model goes beyond the standard approach in the dial-a-ride problem by providing a realistic probabilistic model for each passenger’s expectations—ultimately allowing the provider to tailor journey offers to the passengers that request the service. Our new approach significantly extends our initial work in [11] by enriching the passenger models, and improving the joint scheduling, routing and pricing algorithm to allow for deviations from requested journeys.

B. Modeling Passenger Expectations

In traditional approaches to the dial-a-ride problem, as well as taxi routing and scheduling, a passenger is simply a request; that is, pick-up and drop-off times, and locations. By viewing a passenger in this way it is not possible to optimize the price for each passenger. This is due to the fact that the service provider must also account for passenger preferences; in other words, how likely a passenger is to accept a journey.

To enrich the passenger model, we account for the probability that a passenger will accept an offer, on top of her request. Two key factors are considered: the price of the offer; and the journey deviation. The journey deviation corresponds to the difference between the time that passengers request pick-up and drop-off, and the actual times. This passenger model forms the basis for price optimization in our new mechanism for scheduling, routing, and pricing.

C. A Market Mechanism for the New Thread

At its heart, the new thread of on-demand transport services is multiple, independent, passengers ordering transportation from a service provider with multiple vehicles. A natural way forward is to use a market, as markets, by definition, exchange goods or services for money.

In this paper, we propose a market mechanism for the new thread of on-demand transport services—profit-motivated service providers with coordinated drivers. Our market mechanism is designed to jointly route and schedule vehicles, and price passengers. In particular, we

introduce a four-stage mechanism, initiated by the service provider generating an offer for passengers, and ending with passengers making a final decision of whether to accept or reject.

It is worth pointing out the key difference between our market mechanism and other negotiation mechanisms used for on-demand transport services. That is, our mechanism ultimately schedules, routes, and *prices* passengers. This is not the case in other approaches. The reason is that the other negotiation mechanisms are used for, in a decentralized fashion, scheduling and routing vehicles. However, there is no consideration of how to optimally price passengers. On the other hand, we are able to optimize passenger pricing (in addition to routing and scheduling vehicles) due to our enriched passenger model, which captures passenger expectations.

D. Key Contributions

In this paper, we introduce a new thread of on-demand transport services—profit-motivated providers and coordinated drivers—and a market mechanism to route, schedule, and price passengers. We summarize our key contributions as follows:

- 1) **Agent-based passenger modeling:** We introduce new models for passengers that are enriched to include expectations for price and deviations from requests, on top of the standard request-based model.
- 2) **Market mechanism:** We propose a new market-mechanism to jointly schedule, route, and price passengers. Offers are generated for passengers via an expected profit maximization algorithm. We also analyze the effect of varying the time between when the mechanism is run, which leads to closed-form analytical expressions that provide key design insights.
- 3) **Simulation results:** The on-demand transport services considered in this paper are a new category and as such there is a genuine need for a business case. To this end, we perform a simulation study to evaluate the potential for profit, and also the efficient service of passengers. We demonstrate that incorporating passenger expectations in our new passenger models can in fact improve profitability over standard fixed price-rate approaches, while maintaining comparable efficiency.

II. MODELING AGENTS

Consider the network consisting of a single on-demand transport service provider and N passengers. The service provider owns a fleet of K *unit capacity* vehicles that all start and finish their journeys at a common depot; each vehicle traveling with average velocity ν . Each passenger has requested pick-up and drop-off locations, which are represented by elements from the set of vertices V in a directed graph G . The directed graph G represents the underlying road network. As such, the set of edges E in G represent direct routes between locations in V .

Associated to each edge $e \in E$ (corresponding to direct routes between locations) are:

- 1) a start location $u \in V$;
- 2) an end location $w \in V$;
- 3) a cost $c_e \in [0, \infty)$ to the service provider to traverse edge $e \in E$;
- 4) and an edge traversal time $\tau_e \in \mathbb{Z}_+$.

The edge cost c_e and the edge traversal time τ_e are found during pre-processing where the service provider solves the shortest path problem between u and v on the underlying road network. Note that when edge e connects the vertices u and w , the traversal time is denoted by $\tau_{u,w}$. This model is appropriate when drivers are salaried or pay a fixed commission, which can be easily incorporated into the edge cost c_e .

So far, we have simply described the basic service provider and passenger model used to model dial-a-ride services. In order to enrich the model for the new thread on-demand services, we need to introduce two new features to the model:

- 1) passengers capable of making a decision whether or not to accept a journey offer;
- 2) and a service provider capable of evaluating the probability a passenger will accept her offer.

A. Passengers

The first step in modeling passengers is descriptive. In particular, we identify two stylized facts, which we believe hold for passengers using on-demand transport services. The concept of stylized facts has been widely used in computational macroeconomics as a means of validating

descriptive models of real markets [12]. Based on the stylized facts, we develop a passenger policy, which is used to model how passenger decide whether or not to accept a journey.

Before introducing the stylized facts and the passenger policy, we define the parameters that determine passenger behavior. First, a request from passenger i consists of:

- 1) a pick-up location $v_{i,p} \in V$;
- 2) a drop-off location $v_{i,d} \in V$;
- 3) a pick-up time interval $(a_i, b_i) \in \{0, 1, 2, \dots\} \times \{0, 1, 2, \dots\}$, with $a_i \leq b_i$;
- 4) and a latest drop-off time $l_i \in \{0, 1, 2, \dots\}$, with $l_i > b_i$.

We note that we only explicitly consider the latest drop-off time (instead of a drop-off time interval) as the earliest pick-up time determines the earliest drop-off time.

In response to the requests, the service provider offers a journey to each passenger, which consists of two components: the deviation of the journey from the request; and the price of the journey. More precisely, the deviations are defined as follows.

Definition 1. Let T_i be the actual pick-up time and L_i be the actual drop-off time, for passenger i . The pick-up interval deviation, denoted by $\gamma_{p,i}$, is defined as

$$\gamma_{p,i} = \begin{cases} a_i - T_i, & \text{if } T_i < a_i \\ T_i - b_i, & \text{if } T_i > b_i \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Similarly, the drop-off time deviation, denoted by $\gamma_{d,i}$, is defined as

$$\gamma_{d,i} = \begin{cases} L_i - l_i, & \text{if } L_i > l_i \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

The deviation is then given by $\delta_i = \gamma_{p,i} + \gamma_{d,i}$.

The next step towards the passenger policy is to introduce two stylized facts for passenger behavior. Stylized facts—widely used in computational economics [13]—are important as they provide a means of justifying the passenger policy as a plausible description of real passenger decision-making relevant to on-demand transport.

Our first stylized fact is as follows.

Stylized Fact 1. *The maximum price a given passenger will pay for a journey and her maximum deviation δ do not vary significantly for passengers that regularly use on-demand transport services. In other words, these parameters will in general vary from passenger to passenger, but not for the same passenger.*

We can justify this stylized fact by observing that passengers change transportation habits with great difficulty [14]. As such, we expect that regular on-demand users will have a well-defined maximum price they are prepared to pay.

Next, we have the second stylized fact.

Stylized Fact 2. *The probability that a passenger will accept an offer decreases when either the price increases with the deviation fixed, or the deviation increases with the price fixed.*

More colloquially, this stylized fact is a formal statement of the intuitive notion that if there is a better deal, more passengers will accept.

Based on the first stylized fact, we propose the following descriptive model for passenger policies used to determine whether or not to accept an offer:

- 1) If $r < r_{\max}$ and $\delta < \delta_{\max}$, then the passenger will accept;
- 2) Otherwise, then the passenger will reject.

Observe that this policy ensures that the first stylized fact holds. Moreover, our descriptive model does not violate the second stylized fact, as the policy is concerned with individual passenger decisions and the stylized fact is concerned with the aggregate. We also note that in principle δ_{\max} can depend on r_{\max} .

B. Service Provider

The aim of the service provider is to schedule, route and price passengers to maximize its expected profit. In the literature, the price-rate has been commonly assumed to be fixed, and only the scheduling and routing were optimized. This means that there is no uncertainty and the expected profit can be maximized by minimizing the cost. The key new feature that we introduce for the service provider is a probabilistic model of each passenger. As we show in Section III, this extra information improves the expected profit over standard approaches.

We now detail the passenger model that the service provider uses. Importantly, we also show that it satisfies both the stylized facts, which means that it is a plausible probabilistic descriptive model of passengers.

In order for the service provider to infer demand at a given price, it requires the probability each passenger will accept her offer. It is necessary to consider the *probability* an offer is accepted as the service provider does not perfectly know the maximum deviations and price that any given passenger will accept. In particular, the probability passenger i accepts her offer is given by

$$\Pr(i \text{ accept}) = \Pr(\delta_i \leq \delta_{i,\max}, r_i \leq r_{i,\max}). \quad (3)$$

Remark 1. *We emphasize that the realizations of the maximum deviation $\delta_{i,\max}$ and the maximum price rate $r_{i,\max}$ are not known to the service provider, only to passenger i .*

To obtain the probability that any given passenger accepts, we assume that the service provider knows the joint probability density function $f(\delta_{i,\max}, r_{i,\max})$. Our assumption that the service provider has statistical knowledge of $(\delta_{i,\max}, r_{i,\max})$ ensures that each passenger is not always charged at the maximum possible price she is prepared to pay—realistic in competitive profit-motivated on-demand transport services. On the other hand, statistical knowledge is enough to enable the service provider to optimize the expected profit, as we detail in Section III.

The density function $f(r_{i,\max}, \delta_{i,\max})$ will typically depend on factors such as time of day, or the location of the service region. We focus on scenarios where the maximum deviations and price-rate are independent, which occur when the factors determining the maximum deviations and the financial factors affecting the price-rate are unrelated. An example is the taxi spot market, where only a small deviation is acceptable. In these scenarios, the density function is separable; i.e.,

$$f(r_{i,\max}, \delta_{i,\max}) = f_r(r_{i,\max})f_\delta(\delta_{i,\max}) \quad (4)$$

We model $f_r(r_{i,\max})$, $f_\delta(\delta_{i,\max})$ via the scaled Beta distribution with parameters (α_r, β_r) , $(\alpha_\delta, \beta_\delta)$ respectively. The reason for this is that the Beta distribution is a flexible distribution, which generalizes a wide variety of distributions with bounded support. The density functions are

given by

$$\begin{aligned} f_r(r_{i,\max}) &= \frac{1}{r_u B(\alpha_r, \beta_r)} \left(\frac{r_{i,\max}}{r_u} \right)^{\alpha_r-1} \left(1 - \frac{r_{i,\max}}{r_u} \right)^{\beta_r-1}, \\ f_\delta(\delta_{i,\max}) &= \frac{1}{\delta_u B(\alpha_\delta, \beta_\delta)} \left(\frac{\delta_{i,\max}}{\delta_u} \right)^{\alpha_\delta-1} \left(1 - \frac{\delta_{i,\max}}{\delta_u} \right)^{\beta_\delta-1}, \end{aligned} \quad (5)$$

where $B(\alpha, \beta)$ is the Beta function, and the densities have support $[0, r_u]$ for $r_{i,\max}$ and $[0, \delta_u]$ for $\delta_{i,\max}$.

It is easy to see that both stylized facts are features of the passenger model that the service provider uses. In particular, as $r_{i,\max}$ and $\delta_{i,\max}$ increase, the probability that a passenger will accept is reduced. This means that both the descriptive passenger models we have detailed will exhibit the features described by the stylized facts.

C. On Prediction

So far, we have argued that our descriptive passenger models satisfy the stylized facts exhibited by real passengers. However, at this point these stylized facts have only been shown to hold under current conditions; that is, the present fixed price-rate approach. The remainder of this paper is concerned with a new routing, scheduling, and pricing approach. Clearly this is a structural change and as such, it is necessary to justify that the stylized facts still hold. This is due to the fact that we are predicting the performance of a socio-technical system under structural changes; known to be notoriously hard to do in many economic settings [13].

Fortunately, the stylized facts (Stylized Facts 1 and 2) are independent of how the service provider performs scheduling, routing, and pricing. This suggests that the stylized facts should hold even when the service provider changes the underlying algorithms and indeed the market. Despite this, it is possible that the parameters of the distributions for the maximum price-rate and deviation may vary when the pricing algorithm is changed. These variations can potentially be overcome by updating the parameters $\alpha_r, \beta_r, \alpha_\delta, \beta_\delta$ to appropriately model the passengers' preferences.

The invariance of the stylized facts to service provider scheduling, routing, and pricing, suggests that key features arising from analysis of our model will be consistent with real-world practice.

III. OUR PROPOSED MARKET MECHANISM

In this section, we propose a new market-based approach for scheduling, routing, and pricing in the new thread of on-demand transport services. We first detail the desiderata that our design should fulfill. We then overview the proposed market mechanism, and detail step-by-step the interactions between the service provider and each passenger.

A. Design Desiderata

The design of on-demand transport services is constrained by the physical (i.e., the vehicle fleet) and financial resources of the service provider, and the expectations of passengers. Ultimately, these constraints determine whether or not the provider is financially viable.

To ensure that the resource constraints of the provider and the expectations of passengers are satisfied, the design of our market mechanism is guided by three key desiderata:

- 1) The service provider should be profitable.
- 2) The passengers that desire the service the most should obtain it.
- 3) There should be a simple interface between passengers and the service provider.

Our first two desiderata are to ensure that the service provider is profitable and that vehicles are allocated to the passengers that most value the service, while the third desiderata brings our design in line with current trends that simplify access to transportation services.

We point out that our second desiderata is closely related to the standard notion of efficiency in mechanism design [15]; that is, an efficient mechanism allocates the service to the passengers that are prepared to pay the most for it. In particular, we can formalize the value of a journey to a given passenger as follows.

Definition 2. Denote the the maximum price for passenger i as $p_{i,\max}$. If the set of serviced passengers is S , then the efficiency of our market mechanism is then defined as

$$\mathcal{E} = \sum_{i \in S} p_{i,\max}. \quad (6)$$

Our third desiderata is to bring our mechanism in line with current trends towards simplifying access to transpotation. The key impact of this desiderata on our approach is that

passengers are not required to price their own journey; instead, the service provider always generates the first offer. This means that passengers do not need to be aware of the behavior of other passengers or how the service provider allocates vehicles. The only decision each passenger needs to make is whether the journey offer is acceptable or not, which in our model is determined by the deviation and price of the offered journey.

B. Overview

The three key design desiderata motivate a market mechanism, where passengers are not required to price their own journeys and the service provider generates offers that maximize its profits. In our approach, we propose the following mechanism structure:

- 1) The service provider makes each passenger an initial offer. The offer to each passenger is based on vehicle allocations and passenger pricing done by the service provider to maximize its average profit.
- 2) Each passenger responds to the initial offer by either rejecting or conditionally accepting the offer. Conditional acceptance is a contract between the passenger and the service provider, which requires the passenger to pay the amount offered unless either the price is raised or the deviation from the requested journey increases.
- 3) The service provider computes the final vehicle journey plans and passenger pricing.
- 4) The passenger either rejects the offer (due to the service provider breaking the contract) or unconditionally accepts, where the passenger is required to pay the provider for the service.

A key aspect of the mechanism is that the service provider generates initial offers for each passenger via optimization of the average profit. There are two sets of variables associated to the optimization problem. First is the set of passengers allocated to the same vehicle, which leads to the allocation $C = \{C_1, \dots, C_K\}$ with C_i corresponding to the passengers allocated to the i -th vehicle. Second is the price-rate, r that each passenger will pay. Formally, the optimization

problem is

$$\begin{aligned}
& \underset{C_1, \dots, C_K, r}{\text{maximize}} && \sum_{S \subseteq N} \left(\sum_{i \in S} r R_i - c_i \right) \prod_{i \in S} \Pr(i \text{ accept}) \prod_{j \in S^c} (1 - \Pr(j \text{ accept})) \\
& \text{subject to} && 0 \leq r \leq r_u,
\end{aligned} \tag{7}$$

where $\Pr(i \text{ accept})$ is given by (3), r is the price-rate offered to each passenger, and c_i is the cost of servicing passenger i .

It is important to note that this optimization problem generalizes the standard formulation for dial-a-ride services. In particular, we optimize over the pricing (encoded in r) and the scheduling and routing (encoded in the sets C_1, \dots, C_K). In contrast, minimum cost scheduling and routing for dial-a-ride services does not consider pricing; i.e., the price-rate r is assumed to be fixed. As such, there is potential for higher expected profit than simply using a minimum cost approach.

We also point out that irrespective of how the scheduling and routing is performed, using our approach in (7) will always lead to a higher expected profit. This is particularly important in the case that it is not tractable to optimally solve the clustering problem and heuristics are required.

While optimally scheduling, routing, and pricing passengers will yield a higher expected profit, it is not straightforward to solve. In particular, the problem in (7) is difficult for two reasons:

- 1) The objective is generally nonlinear and also non-convex.
- 2) The number of sets C_1, \dots, C_K to be searched is even greater than in the standard approaches as there are no feasibility constraints that constrain the passengers able to be serviced by a given vehicle. This occurs because we allow deviations from passenger requests.

To alleviate the two difficulties in solving (7), we adopt the “cluster-then-price” strategy; that is, we solve the problem (7) in two stages. In the first stage we cluster passengers into groups that are all served by the same vehicle. We note that the cluster formation is not a straightforward extension of standard approaches (e.g., [16]) due to the profit-based objective. We detail our clustering algorithm in Section III-C. We then show how the price-rate offered to

each passenger is obtained by pricing the passengers based on the clustering. The remainder of section details the interactions between the service provider and passengers in our mechanism—including how final vehicle allocations and prices are computed.

Our proposed market mechanism is summarized in Fig. 1, where dependencies are illustrated by arrows between each stage.

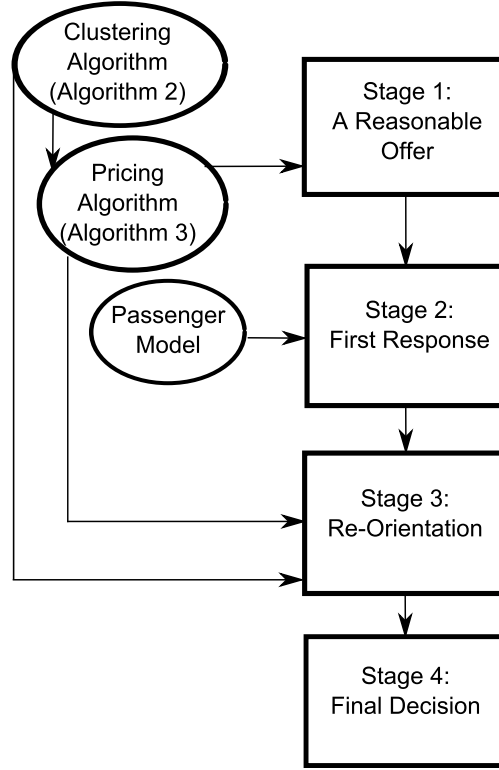


Fig. 1: Overview of the proposed market mechanism for pricing and vehicle allocation in on-demand transport networks.

C. Stage 1: A Reasonable Offer

In the first stage of the mechanism, the service provider generates an offer for each passenger i . The offer consists of:

- 1) maximum deviation, γ_i ;
- 2) and the price, $p_i = rR_i$, where R_i is the distance between pick-up and drop-off locations for passenger i .

To allow deviations from passenger requests during clustering, we introduce a probabilistic feasibility constraint, which is closely related to the probability that a passenger will accept the journey. The constraint is defined as follows.

Definition 3. *A journey is ϵ -feasible for passenger i , if*

$$\Pr(\delta_{\max} \geq \delta_i) \geq 1 - \epsilon, \quad (8)$$

where δ_{\max} is a random variable representing the unknown maximum deviation for passenger i . Moreover, a cluster is said to be ϵ -feasible if each passenger in the cluster is ϵ -feasible, with the set of ϵ -feasible clusters denoted by \mathcal{F}_ϵ .

Intuitively, the notion of ϵ -feasible allocations generalizes the standard notion of feasible clustering. That is, standard hard constraint clustering is 0-feasible; i.e., each passenger must have a feasible journey with probability one. Importantly, ϵ -feasible allocations are a key feature in our approach, as by relaxing the feasible constraints it is possible for the service provider to offer journeys with deviations from requests to passengers. Ultimately, this can allow the provider to service more passengers with a single vehicle, which (as we show in Section V) leads to higher profits and efficiency of the market mechanism.

Passenger clustering: We now develop a clustering algorithm that minimizes the cost of clusters subject to the probabilistic feasibility constraint detailed in Definition 3. In order to form passenger clusters, we determine whether or not the addition of a passenger to a cluster is ϵ -feasible. This is achieved by computing the times that the inserted passenger will be picked up and dropped off. To reduce the complexity of the insertion algorithm, a passenger can be inserted only if the new passenger does not affect the deviations of the passengers already in the cluster, which means that previously clustered passenger pick-up and drop-off times are not changed after the passenger has been allocated to a cluster.

To illustrate, consider a passenger 0 to be inserted into cluster $C_j = \{k_1, \dots, k_{|C_j|}\}$, where passengers in C_j are in the order that they are serviced. There are three situations to consider: insert passenger 0 before passenger k_1 ; insert passenger 0 between passengers k_l and k_{l+1} , with $1 \leq l < |C_j|$; and insert passenger 0 after passenger $k_{|C_j|}$.

To determine whether a passenger can potentially be inserted into the cluster, there are three checks that need to be performed:

- C1:** Can the passenger be inserted without changing the deviation of previously clustered passengers?
- C2:** If C1 holds, what are the pick-up and drop-off times with the minimum deviation (i.e., δ_i)?
- C3:** Is the cost of traveling to the passenger's pick-up location and from the drop-off location less than any other previously checked cluster?

The checks are performed using the insertion algorithm detailed in Algorithm 1. In particular, the algorithm attempts to insert passengers such that the new passenger does not change the deviations of previously clustered passengers—corresponding to **C1**. In the case that it is possible to insert the passenger without affecting other clustered passengers, the insertion algorithm determines whether the passenger is ϵ -feasible and computes the pick-up and drop-off times to minimize the deviation (implementing **C2**). Finally, the cost is computed and compared with previously checked clusters, and the lowest cost, c_{best} , insertion is updated (based on **C3**).

The minimum deviation pick-up and drop-off times (for **C2**) are computed using simple inequality tests due to the fact that the objective is linear; i.e., it is $\delta = \gamma_p + \gamma_d$ (from Section II). In particular, there are three types of insertion tests: before the first passenger in the cluster; between passengers j and $j + 1$; and after the last passenger. To illustrate, consider the potential insertion of passenger 0 between passengers j and $j + 1$. First, the new passenger's journey must fit. This means that $L_j + \tau_{j,0} + \tau_{0,0} + \tau_{0,j+1} < T_{j+1}$. Next, we choose the minimum deviation insertion. This is done by checking whether $T_{j+1} - \tau_{0,j+1} \leq l_0$, which means that the deviation can only be caused by the pick-up time. The pick-up time is then chosen to minimize the deviation. The case where $T_{j+1} - \tau_{0,j+1} > l_0$ can be treated similarly.

The cost is computed as

$$c_{i,ins} = c_{l-1,i} + c_{i,l}, \quad (9)$$

where passenger i is inserted between passenger $l - 1$ (if $l = 0$, then this corresponds to the depot), and passenger l (this may also be the depot).

Algorithm 1 Insertion algorithm.

procedure INSERTION(c_{best}, i, C_j)

Set passenger to be inserted as i .

Compute cost $c_{i,ins}$ for cluster $\{i\}$ using (9).

if $c_{i,ins} < c_{best}$ **then**

$c_{best} = c_{i,ins}; l_{best} = l; j_{best} = j.$

end if

Set cluster $C_j = \{k_1, \dots, k_M\}$, where $M = |C_j|$.

for $1 \leq l \leq |C_j| + 1$ **do**

Perform check **C1**.

Compute pick-up and drop-off times t_p^* and t_d^* (see discussion).

if $\Pr(\delta_{\max} \geq t_p^*, \gamma_{\max} \geq t_d^*) \geq 1 - \epsilon$ **then**

Compute cost $c_{i,ins}$ using (9).

if $c_{i,ins} < c_{best}$ **then**

$c_{best} = c_{i,ins}; l_{best} = l; j_{best} = j.$

end if

end if

end for

return $j_{best}, l_{best}, c_{best}$.

end procedure

With the insertion algorithm in hand, it is now possible to describe our cluster formation algorithm. The cluster formation algorithm searches through the current clusters for each passenger i to determine whether the passenger can be inserted. After checking each potential cluster, passenger i is inserted into the lowest cost and feasible cluster, and the next passenger is considered. Our cluster formation algorithm is detailed in Algorithm 2.

Passenger pricing: With the cluster formation algorithm in hand, we now introduce the joint scheduling, routing, and pricing algorithm used by the service provider to generate expected

Algorithm 2 Cluster formation algorithm with probabilistic feasibility constraints.

procedure CLUSTERFORMATION

Randomly choose a unique index in $\mathcal{N} = \{1, 2, \dots, N\}$ for each passenger.

Initialize cluster $C_1 = \{1\}$;

Initialize allocation $\pi = \{C_1\}$ and partition index $\mathcal{I} = \{1\}$.

Initialize the set of unclustered passengers $\mathcal{U} = \{2, \dots, N\}$.

while $\mathcal{U} \neq \emptyset$ **do**

Set $\mathcal{J} \leftarrow \mathcal{I}$

while $\mathcal{J} \neq \emptyset$ **do**

Randomly choose an element $j \in \mathcal{J}$.

Update $j_{best}, l_{best}, c_{best}$ using $\text{Insertion}(c_{best}, i, C_j)$ (see Algorithm 1).

$\mathcal{J} \leftarrow \mathcal{J} \setminus j$

end while

Update $C_{j_{best}}$ by inserting i before l_{best} ; $\mathcal{U} \leftarrow \mathcal{U} \setminus i$; $\mathcal{J} \leftarrow \emptyset$.

end while

return Cluster allocation π

end procedure

profit maximizing offers. The algorithm approximately solves the following problem:

$$\begin{aligned}
 & \underset{r, \epsilon}{\text{maximize}} && \mathbb{E}[P(r)] \\
 & \text{subject to} && 0 \leq r \leq r_u \\
 & && 0 \leq \epsilon \leq 1
 \end{aligned} \tag{10}$$

where $\mathbb{E}[P(r)]$ is the expected profit at price rate r , which is given by

$$\mathbb{E}[P(r)] = \sum_{S \subset N} \left(\sum_{i \in S} p_i - c_i \right) \prod_{i \in S} \Pr(i \text{ accept}) \prod_{j \in S^c} (1 - \Pr(j \text{ accept})), \tag{11}$$

with $\Pr(i \text{ accept})$ given by (3). Importantly, all price rates greater than r_u are rejected by the passengers with probability one (this follows from (5)). The price offered to each passenger i is

Algorithm 3 Joint pricing and clustering algorithm to maximize the expected profit for Stage 1 of our market mechanism.

procedure EXPECTEDPROFITMAXIMIZATION

Initialize ϵ and set $P_{opt} \leftarrow 0$.

while $\epsilon > 0$ **do**

Set $\epsilon \leftarrow \epsilon - \epsilon_{step}$.

Solve (10) to obtain the optimal expected profit P^* , with ϵ fixed.

if $P^* > P_{opt}$ **then**

Set $P_{opt} \leftarrow P^*$; $r_{opt} \leftarrow r$.

Set $C_{opt} \leftarrow \{C_1, C_2, \dots\}$ (corresponding to clusters from Algorithm 2).

end if

end while

return $P_{opt}, r_{opt}, C_{opt}$.

end procedure

then given by $p_i = rR_i$, where R_i is the distance from passenger i 's pick-up to drop-off.

Our solution clusters passengers based on the probabilistic feasibility constraint with parameter $\epsilon_1 \approx 1$ and then optimizing the price using a standard scalar nonlinear optimization algorithm (i.e., a descent algorithm). This is repeated for parameter $\epsilon_{k+1} = \epsilon_k - \epsilon_{step}$ until $\epsilon_{k+1} < 0$, with $0 < \epsilon_{step} \leq 1$. The pricing and clustering solution that maximizes the expected profit over probabilistic feasibility parameters $\epsilon_1, \epsilon_2, \dots$ is then chosen. As such, at the end of Stage 1, the price rate r and passenger clusters C_1, C_2, \dots are obtained by the service provider. The procedure is summarized in Algorithm 3.

The final output of the expected profit maximization algorithm is an offer consisting of a journey deviation and price for each passenger. This offer is then communicated to the passenger, and the service provider waits for the passengers response, given in the next stage.

D. Stage 2: First Response

In the second stage of the negotiation, each passenger makes a preliminary decision to accept or reject the *conditional journey* offered by the service provider. If passenger i accepts, it means that she has accepted the journey offer as long as the service provider does not change the offer. On the other hand, if passenger i rejects the offer then she is no longer interested in a journey with the service provider.

We emphasize that if the passenger accepts, then she has agreed to a contract with the service provider; that is, she must pay for the service unless the service provider either raises the price or increases the journey deviation. We note that this type of contract is standard for other transportation services, such as pre-booked trains or buses.

E. Stage 3: Re-Orientation

In the third stage of the negotiation, the service provider has additional information. In particular, the service provider knows both what the users that conditionally accept are prepared to pay, and which users have rejected the offer.

As not all passengers will usually accept their offer, the service provider must update the passenger clusters. Although the cluster sizes will change if not all passengers accept, no passengers are allocated to different clusters. This ensures that the maximum journey deviation for each passenger does not change.

To obtain final prices for each passenger, the service provider solves the maximin profit problem over the passengers that have accepted in the previous stage. Let S^* be the passengers that accepted their offers in the previous stage and Q the final profit obtained from the market mechanism. The maximin profit problem is then

$$\max_{\{r_i\}_{i \in S^*}} \min Q. \quad (12)$$

We note that in this stage of the mechanism, different passengers can be offered journeys at different price rates.

As the previous stages of the mechanism have revealed a lower bound on the maximum price each user is prepared to pay, the maximin profit problem in (12) is equivalent to finding

the set of passengers

$$S_c^* = \arg \max_{S_c \subset S^*} \sum_{k \in S_c} rR_k - c_{S_c}, \quad (13)$$

and then pricing passengers such that the passengers in S_c^* are charged rR_k , $k \in S_c^*$, while the other passengers are charged at a higher price rate. This allows these passengers to find another service or for the provider to subcontract the journeys, which avoids losses.

The service provider uses the pricing strategy in Stage 3 to ensure that the passengers in desirable clusters accept, which in turn maximizes the service provider's profit. That is, the service provider will raise the price of passengers that have conditionally accepted so that they do not require the provider to be exposed to large losses. Although undesirable from the perspective of service provider reputation, we believe that this strategy is likely to be necessary in real-world practice. This is due to the fact that service providers have both physical (i.e., fleet size) and financial (i.e., initial capital) constraints. As such, it is not possible to service all passengers that might accept without either investing in a larger fleet size or hiring additional vehicles. To cope with these additional costs, it is necessary for passengers to be charged more when they are difficult to serve.

F. Stage 4: Final Decision

In the fourth stage of the negotiation, the remaining passengers make their final decision based on the latest offer from the service provider. Each passenger either *unconditionally accepts* the final offer, or *rejects* it; i.e, passengers that conditionally accepted in Stage 2 now either accept the offer or reject otherwise. We emphasize that a passenger cannot freely reject the offer unless the provider has increased the price; otherwise, the passenger must pay for the journey.

We point out that the service provider cannot always service all passengers that accept; either because there are not enough vehicles, or the passenger would cause a net loss for the provider. As such, it is highly desirable from the perspective of financial solvency of the provider to be able to raise the price and allow the passenger to find alternative transport. If the passenger still accepts the journey even after the price has been raised, then the service provider can use the additional revenue to hire an additional vehicle to service the passenger.

At the end of Stage 4 of the mechanism, all passengers to be serviced are known to the service provider, are priced, and have been allocated to vehicles. Moreover, each vehicle has a journey plan. The performance of our market mechanism is evaluated in Section V. In the next section, we analyze the role of the time interval between mechanism runs, which determines the mechanism rate.

IV. MECHANISM PARAMETER DESIGN

A key assumption in our market mechanism is that the passengers are known to the provider before the beginning of the mechanism. For dial-a-ride services this is known as the static scenario, and is often problematic as passengers can make requests after one run of the mechanism and before the next. This was solved by allowing dynamic arrivals, where passengers can be inserted while vehicles are on the road. However, the dynamic approach cannot be directly used with our mechanism without statistics for the locations and the prices dynamic passengers would be prepared to pay. While such an approach is in principle possible using historical passenger prices, and pick-up and drop-off location data, an unprecedented level of data refinement would be required.

To resolve this issue, we instead adapt the rate our market mechanism is run. Importantly, the mechanism rate is in fact a fundamental feature of any on-demand transport market mechanism. There are two key parameters that determine the mechanism rate: the probability that a passenger request is ignored; and the probability that a passenger cannot be serviced before the next mechanism run.

In this section, we derive simple analytical expressions for the probability a request is ignored, P_{ignore} , and the probability a passenger cannot be serviced in time, $P_{overtime}$. The key purpose of the analytical expressions are to guide design of the mechanism rate. In particular, we demonstrate the tradeoff between P_{ignore} and $P_{overtime}$, as the time between mechanism runs is increased.

A. Analysis

Our analysis of P_{ignore} and $P_{overtime}$ is based on a simplified probabilistic model of passengers and vehicles. Although the simplifications lead to a coarse approximation of real-world on-demand transport networks, conclusions from our analysis are supported by intuitive explanations. The key assumptions are as follows:

- 1) The time between mechanism runs is T minutes and the corresponding rate is $R = 1/T$.
- 2) Each vehicle services only a single passenger per interval between mechanism runs. This is reasonable for short intervals T .
- 3) The time for a vehicle to travel a distance z is given by νz , where ν is the average velocity.
- 4) The time each request arrives forms a homogeneous Poisson process with rate κ . Moreover, the time between request delivery and the desired pick-up time, Δ , is exponentially distributed with mean $1/\lambda$.
- 5) Pick-up and drop-off locations are distributed according to a Poisson point process with intensity ζ . Moreover, the drop-off location is the closest point to the pick-up locations, which means that the distribution of the distance is²

$$f_Z(z) = e^{-\zeta\pi z^2} 2\pi\zeta z. \quad (14)$$

We now turn to analysis of P_{ignore} and $P_{overtime}$. Our analysis is based on two new analytical expressions for the probabilities. First, the probability a request is ignored, P_{ignore} , is given in the following proposition.

Proposition 1. *The probability a request is ignored is given by*

$$P_{ignore} = 1 - \frac{1}{\lambda T} (1 - e^{-\lambda T}). \quad (15)$$

Proof: See Appendix A. ■

Observe that as $T \rightarrow \infty$, $P_{ignore} \rightarrow 1$, which means that as the interval between mechanism runs increases, the probability a request is ignored tends to one, which is consistent with intuition. On the other hand, as $T \rightarrow 0$, $P_{ignore} \rightarrow 0$.

²This result follows immediately by considering the probability that there is no point within radius R from the origin, which is given by $e^{-\pi R^2}$.

The probability that a passenger cannot be serviced before the next mechanism run, $P_{overtime}$, is given in the following proposition.

Proposition 2. *The probability a passenger cannot be serviced before the next mechanism run is given by*

$$P_{overtime} = \left(1 - \frac{2\pi\zeta T}{\nu}\right) e^{-\pi\zeta T^2/\nu^2} + \frac{\nu}{T\sqrt{\zeta}} \left(\frac{1}{2} - Q\left(\frac{T\sqrt{2\pi\zeta}}{\nu}\right)\right) \quad (16)$$

where $Q(\cdot)$ is the Q -function, a standard special function defined in (22).

Proof: See Appendix B. ■

Observe that as $T \rightarrow \infty$, $P_{overtime} \rightarrow 0$, and as $T \rightarrow 0$, $P_{overtime} \rightarrow 1$. This is an intuitive result as it simply states that when the interval is large the probability that a passenger cannot be serviced in time approaches zero. Importantly, this observation means that there is a tradeoff between P_{ignore} and $P_{overtime}$. In other words, it is not possible to avoid passengers being ignored and not being serviced, with probability one using the same mechanism rate.

B. Tradeoff

It is not immediately obvious from Propositions 1 and 2, which factors are key in determining the tradeoff between P_{ignore} and $P_{overtime}$. As such, we now examine the behavior of P_{ignore} and $P_{overtime}$ numerically.

Fig. 2 plots the tradeoff between P_{ignore} and $P_{overtime}$ for varying λ and ζ , based on our analysis in Section IV-A. Observe that increasing λ , also leads to an increase in P_{ignore} . On the other hand, an increase in ζ leads to increase in $P_{overtime}$. As such, the intersection between P_{ignore} and $P_{overtime}$ (the crossover rate) reduces when λ or ζ are increased.

V. SIMULATION RESULTS

With our mechanism for the new thread of on-demand transport services in hand, we are ready to present an initial business case for our mechanism. The benchmark we use for comparison is the fixed price-rate approach, where passengers are charged the same price-rate, irrespective of the number of potential passengers or the requested journeys. In particular, under




Fig. 2: Plot of the tradeoff between P_{ignore} (from Proposition 1) and P_{overtime} (from Proposition 2), for varying time between mechanism runs T . We assume vehicles travel at 20 km/hr.

the fixed price-rate policy, each passenger is charged at a rate given by the expected maximum price-rate passengers are prepared to pay, given by $\mathbb{E}[r_{i,\max}]$, which can be easily obtained from (5).

The remainder of this section consists of two parts. First, we compare our mechanism as detailed in Section III with the standard fixed price-rate policy. The performance of our mechanism is evaluated in terms of both the expected profit and the expected efficiency (based on Definition 2). We demonstrate key trends in these performance metrics with variations in the number of potential passengers, passenger demand, and the clustering parameter ϵ (see Definition 3). Second, we evaluate the performance of a modification of our mechanism in a scaled-up system with up to 100 passengers. We describe this modification and demonstrate the performance as the number of potential passengers are varied, in comparison with the fixed price-rate approach.

A. Key Trends

We now illustrate key trends in the expected profit and expected efficiency, based on a network setup with $K = 5$ drivers and up to $N = 13$ potential passengers. Importantly, this setting is practical—despite the small scale—when the serviced region is partitioned and separate negotiations are performed in each partition. This is one approach that reduces the need for suboptimal heuristics for pricing, which can ensure practical negotiation run-times. We evaluate an alternative approach based on a heuristic approach on a scaled-up network in Section V-B.

The pick-up and drop-off locations of passengers in the network are drawn from real locations in Prague, Czech Republic: $K = 5$ drivers; average vehicle velocity $\nu = 30$ km/hr; cost/km of 0.4 euros/km; and a maximum price-rate for each passenger of 3 euros/km. We further assume that there is an hour between each mechanism run, which means that the beginning of a passengers pick-up interval is uniformly distributed over the 60 minute interval. The maximum length of each passenger's pick-up interval is 10 minutes, with the actual interval length uniformly distributed.

We first demonstrate the performance mechanism with clustering using hard constraints (i.e., we set $\epsilon = 0$ and $\epsilon_{step} > 1$). We note that the mechanism with a fixed price-rate given by $\mathbb{E}[r_{i,\max}]$ and clustering with hard constraints is used as a benchmark.

Fig. 3 plots the expected profit per mechanism run for a varying number of potential passengers, N . The key observation is that our mechanism with optimized price-rate always improves the profit over the fixed price-rate approach, although the improvement depends on the passenger demand; i.e., how much passengers are willing to pay. Observe that there is a significant improvement in expected profit for both low and high demand, corresponding to $\alpha_r = 1, \beta_r = 3$ and $\alpha_r = 1, \beta_r = 1$, respectively. We also point out that the rate of increase in the expected profit with the number of passengers, N , is reduced as N increases. This is due to a saturation effect, where increasing the number of potential passengers in the network does not significantly improve the profit.

Next, we turn to the expected efficiency (see (2)) of our mechanism demonstrated in Fig. 4.




Fig. 3: Plot of the expected profit for our mechanism with optimized price-rate and the standard fixed price-rate approach, for varying number of potential passengers, N . High demand corresponds to $\alpha_r = 3$, $\beta_r = 1$, medium demand corresponds to $\alpha_r = \beta_r = 1$, and low demand corresponds to $\alpha_r = 1$, $\beta_r = 3$. Simulation parameters: $K = 5$; average velocity of 30 km/hr; cost/km of 0.4; maximum price-rate of 3 euros/km; and hard feasibility constraints are enforced.

In this case, observe that our mechanism improves the expected profit over the fixed price-rate approach for low demand with $N \geq 7$. This means that our mechanism outperforms the fixed price-rate approach in both expected profit and efficiency. On the other hand, the fixed price-rate approach has a higher efficiency for high demand.

Next, we demonstrate the performance improvements that can be obtained by using our optimized clustering algorithm, where $\epsilon_{step} \leq 1$. In particular, this improves the efficiency in the case of high demand.

In Table I, the expected profit is compared with the number of potential passengers, with and without optimized clustering. The optimized clustering solves (10) with $\epsilon_{step} = 0.2$, while the algorithm without optimized clustering enforces the hard constraints (i.e., $\epsilon = 0$). In both




Fig. 4: Plot of the expected efficiency for our mechanism with optimized price-rate and the standard fixed price-rate approach, for varying number of potential passengers, N . High demand corresponds to $\alpha_r = 3$, $\beta_r = 1$, medium demand corresponds to $\alpha_r = \beta_r = 1$, and low demand corresponds to $\alpha_r = 1$, $\beta_r = 3$. Simulation parameters: $K = 5$; average velocity of 30 km/hr; cost/km of 0.4 euros/km; maximum price-rate of 3 euros/km; and hard feasibility constraints are enforced.

cases, the price-rate is optimized. Observe that our optimized clustering algorithm improves the expected profit, even with optimal pricing. In particular, gains of up to 4 euros per negotiation can be achieved. This means that the expected profit for the fixed price-rate approach can be improved significantly (up to 10 euros per negotiation with $\alpha_r = \beta_r = 1$) by using both optimized pricing and clustering.

In Table II, the expected efficiency (in euros/m) is compared with the number of potential passengers, with and without optimized clustering. Observe that the expected efficiency of optimized clustering outperforms the expected efficiency without optimized clustering. Importantly, optimizing the clustering ensures that the expected efficiency of our mechanism is comparable with the efficiency using the fixed price-rate approach. As such, our mechanism can outperform

TABLE I: Expected profit per mechanism run (in euros) with and without optimized clustering.

Parameters: $\alpha_r = \beta_r = 1$, $\alpha_\delta = 3$, $\beta_\delta = 1$.

Potential Passengers N	5	7	9	11	13
With Optimized Clustering (euros)	17.5	23.2	26.3	28.1	30.0
Without Optimized Clustering (euros)	14.5	18.8	22.0	25.5	27.8

the fixed price-rate approach in both expected profit and expected efficiency.

TABLE II: Expected efficiency per mechanism run (in euros) with and without optimized clustering. Parameters: $\alpha_r = \beta_r = 1$, $\alpha_\delta = 3$, $\beta_\delta = 1$.

Potential Passengers N	5	7	9	11	13
Optimized Clustering (euros/m)	0.0038	0.0051	0.0063	0.0074	0.0085
Without Optimized Clustering (euros/m)	0.0030	0.0038	0.0044	0.0052	0.0057

B. The Effect of Network Scaling

We now turn to simulation results for a scaled-up network with 30 vehicles and up to 100 passengers. The average velocity, cost/km, maximum price-rate, and pick-up interval distributions are the same as in Section V-A.

In order to apply our mechanism to a network of this scale, it is necessary to modify our mechanism by approximating the objective of the optimization problem in (10) by truncating the sum involved in the expectation in (10). Moreover, clustering is limited to the finding the first feasible cluster instead of the full search detailed in Algorithm 1. Using these heuristics, next we compare our mechanism with the fixed price-rate approach without clustering; i.e., each vehicle serves up to one passenger per negotiation run.

Table III demonstrates the effect of the heuristic on the performance of the network. Observe that in all cases the optimized pricing in our mechanism outperformed the fixed price-rate approach, in terms of expected profit. Note that the expected profit for the fixed price-rate

approach is approximately constant, irrespective of the number of passengers. This is due to the saturation effect also observed in Fig. 3. On the other hand, the expected profit for the optimized pricing approach reduces as the number of potential passengers increases, in contrast with the results in Section V-A. This is due to the approximation of the objective, which is less accurate as the number of potential passengers increases.

TABLE III: Expected profit per mechanism run (in euros). Parameters: $\alpha_r = \beta_r = 1$, $\alpha_\delta = 3$, $\beta_\delta = 1$.

Potential Passengers N	30	60	90
Optimized Pricing (euros)	83.7	66.7	52.9
Fixed Price-Rate (euros)	52.6	51.7	52.4

C. Summary Of Simulation Results

We summarize the key results from our simulation study as follows.

- 1) Our mechanism outperforms the fixed price-rate approach in terms of expected profit, even without optimized clustering (i.e., $\epsilon = 0$ and $\epsilon_{step} > 1$) or using heuristic approximations.
- 2) Using optimized clustering with $\epsilon_{step} \leq 1$ both increases the expected profit and the expected efficiency of our mechanism. Moreover, the expected profit is significantly higher than the fixed price-rate approach, irrespective of the passenger demand, and also has comparable efficiency.
- 3) Increasing the passenger demand (i.e., varying the parameters α, β in the distribution) leads to an increase in the expected profit.
- 4) Increasing the number of potential passengers using our optimized clustering algorithm increases the expected profit; however, the expected profit saturates at a large number of potential passengers.
- 5) Using the heuristic modification of our mechanism for scaled-up networks (described in Section V-B) outperforms the fixed price-rate approach in terms of expected profit.

However, the expected profit reduces as the number of potential passengers is increased. This suggests that an appropriate method to deal with large-scale networks is to partition the network by simultaneously running several small-scale negotiations. In this case, approximations of the objective function are not required, which in turn ensures that the expected profit can increase as the number of potential passengers increases.

VI. CONCLUSIONS

We have proposed a new market mechanism for a new thread of on-demand transport services, which enables negotiations with passengers to both increase provider profits and select passengers that value the service the most. A key feature of our mechanism is that it jointly optimizes scheduling, routing and passenger pricing; in sharp contrast with standard approaches for services targeted at the elderly and disabled, and taxis.

Our mechanism is based on a new agent-based model, which emphasizes the role of provider profit in allocating resources. In particular, we have developed new models that incorporate price-based preferences for both the passengers and the provider. We also consider the effect of deviations from passenger requests in our resource allocation.

Our simulations results for new thread of on-demand transport services demonstrated that our market mechanism improves the profitability of the service provider, compared with standard fixed price-rate approaches, while maintaining comparable efficiency.

ACKNOWLEDGMENTS

Access to computing and storage facilities owned by parties and projects contributing to the National Grid Infrastructure MetaCentrum, provided under the programme "Projects of Large Infrastructure for Research, Development, and Innovations" (LM2010005), is greatly appreciated.

APPENDIX A

PROOF OF PROPOSITION

Consider a time interval $[nT, (n+1)T]$ ($n \in \mathbb{N}$) between two consecutive mechanism runs, with $a \in [nT, (n+1)T)$. Then, the probability that a passenger's request is ignored is given by

$$P_{\text{ignore}} = 1 - \Pr(a + \Delta \geq (n+1)T). \quad (17)$$

Since request arrive according to a homogeneous Poisson process, the following lemma gives the distribution of the earliest pick-up time, a .

Lemma 1. *The earliest possible pick-up time a is uniformly distributed in $[nT, (n+1)T)$, conditioned on the number of requests in the interval.*

This means that for a random request in the interval $[nT, (n+1)T)$, a is uniformly distributed.

Using Lemma 1 and conditioning on the time between the request arrival and desired pick-up time, Δ , we have

$$P_{ignore} = \int_0^\infty \Pr(a \geq (n+1)T - \delta | \Delta = \delta) \lambda e^{-\lambda\delta} d\delta. \quad (18)$$

Integrating by parts, we obtain the required result.

APPENDIX B

PROOF OF PROPOSITION

As for Proposition A, consider a time interval $[nT, (n+1)T)$ ($n \in \mathbb{N}$) between two consecutive mechanism runs, with $a \in [nT, (n+1)T)$. Then, the probability that a passenger cannot be serviced before the next mechanism run is given by

$$\begin{aligned} P_{overtime} &= \Pr(a + \nu z \geq (n+1)T) \\ &= \int_0^\infty \Pr(a \geq (n+1)T - \nu z) f_Z(z) dz, \end{aligned} \quad (19)$$

which follows by conditioning on the distance, Z . We note there is no loss in generality by assuming that the pick-up location is at the origin by Slivnyak's theorem [17].

Using Lemma 1 and the fact that $a \in [nT, (n+1)T)$, it follows that

$$\begin{aligned} P_{overtime} &= \frac{\nu}{T} \int_0^{T/\nu} z f_Z(z) dz + \int_{T/\nu}^\infty f_Z(z) dz \\ &= \frac{\nu}{T} \int_0^{T/\nu} 2\pi\zeta z^2 e^{-\pi\zeta z^2} dz + e^{-\pi\zeta T^2/\nu^2}. \end{aligned} \quad (20)$$

Next we integrate by parts³ to obtain

$$P_{overtime} = \frac{\nu}{T} \left(\frac{T}{\nu} e^{-\pi\zeta T^2/\nu^2} + \int_0^{T/\nu} e^{-\pi\zeta z^2} dz \right) + e^{-\pi\zeta T^2/\nu^2}. \quad (21)$$

³We use the fact that $\int 2\pi\zeta z e^{-\pi\zeta z^2} dz = -e^{-\pi\zeta z^2}$.

We then apply the change of variables $u = \sqrt{2\pi\zeta}z$ and the definition of the Q -function, which is given by

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-u^2/2} du. \quad (22)$$

Using the fact that $Q(0) = \frac{1}{2}$, we then have

$$P_{overtime} = \left(1 - \frac{2\pi\zeta T}{\nu}\right) e^{-\pi\zeta T^2/\nu^2} + \frac{\nu}{T\sqrt{\zeta}} \left(\frac{1}{2} - Q\left(\frac{T\sqrt{2\pi\zeta}}{\nu}\right)\right), \quad (23)$$

as required.

REFERENCES

- [1] J.-F. Cordeau, "The dial-a-ride problem: models and algorithms," *Annals of Operations Research*, vol. 153, pp. 29–46, 2007.
- [2] D. Barbucha, "A multi-agent approach to the dynamic vehicle routing problem with time windows," in *Proc. of the International Conference on Computational Collective Intelligence*, 2013.
- [3] C. Cubillos, F. Guido-Polanco, and C. Demartini, "Madarp: multi-agent architecture for passenger transportation systems," in *Proc. IEEE International Conference on Intelligent Transport Systems*, 2005.
- [4] H. Yang, C. Fung, K. Wong, and S. Wong, "Nonlinear pricing of taxi services," *Transportation Research Part A: Policy and Practice*, vol. 44, no. 5, pp. 337–348, 2010.
- [5] K. Seow, N. Dang, and D.-H. Lee, "A collaborative multiagent taxi-dispatch system," *IEEE Transactions on Automation Science and Engineering*, vol. 7, no. 3, 2010.
- [6] R. Bai, J. Li, J. Atkin, and G. Kendell, "A novel approach to independent taxi scheduling problem based on stable matching," *Journal of the Operational Research Society*, vol. 65, 2014.
- [7] H. Yang and T. Yang, "Equilibrium properties of taxi markets with search frictions," *Transportation Research Part B: Methodological*, vol. 45, no. 4, pp. 696–713, 2011.
- [8] R. Balan, N. Khoa, and L. Jiang, "Real-time trip information service for a large taxi fleet," in *Proc. International Conference on Mobile Systems, Applications, and Services*, 2011.
- [9] A. Glaschenko, A. Ivaschencko, G. Rzevski, and P. Skobolev, "Multi-agent real time scheduling system for taxi companies," in *Proc. International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2009.
- [10] C. Zeng and N. Oren, "Dynamic taxi pricing," in *In Proc. European Conference on Artificial Intelligence*, 2014.
- [11] M. Egan and M. Jakob, "A profit-aware negotiation mechanism for on-demand transport services," in *In Proc. European Conference on Artificial Intelligence*, 2014.
- [12] R. Cont, "Empirical properties of asset returns: stylized facts and statistical issues," *Quantitative Finance*, vol. 1, pp. 223–236, 2001.
- [13] L. Tesfatsion, "Introduction to the special issue on agent-based computational economics," *Journal of Economic Dynamics and Control*, vol. 25, no. 3–4, pp. 281–292, 2001.
- [14] B. Verplanken and W. Wood, "Interventions to break and create customer habits," *Journal of Public Policy and Marketing*, vol. 25, no. 1, pp. 90–103, 2006.

- [15] Y. Shoham and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009.
- [16] Y. Dumas, J. Desrosiers, and F. Soumis, "The pickup and delivery problem with time windows," *European Journal of Operations Research*, vol. 54, pp. 7–22, 1991.
- [17] F. Baccelli and B. Blaszczyzyn, *Stochastic Geometry and Wireless Networks: Volume 1: Theory*. Now Publishers, Inc., 2009.